

Geometric statistical inference

Vipul Periwal

Department of Physics, Princeton University, Princeton, New Jersey 08544

Finite sample size corrections to the reparametrization-invariant solution of the inverse problem of probability are computed, and shown to converge uniformly to the correct distribution.

I. INTRODUCTION

A basic problem in statistics, with applications in divers fields, is the determination of the probability distribution that underlies a finite set of experimental results involving continuous variables. From the practical point of view, one usually has a definite finite-parameter model for the probability distribution on theoretical grounds, and the parameters are fixed by fitting to the data set. As is obvious, this parametric inference introduces a theoretical bias, since the true distribution may not even lie in the parameter set chosen. It is therefore of interest to attempt a direct determination of the probability distribution without resorting to finite-dimensional approximations. Of course, for any finite data set, one obtains only a probabilistic description of the probability distribution, but the spread of possible probability distributions decreases as the size of the data set is increased [1].

This set-up is fundamental in pattern theory, a context in which reparametrization invariance is of great importance. In visual information processing algorithms, for example, if the inferred pattern depends on, say, the orientation of the object, then the object recognition capabilities of the algorithm are not likely to be efficient. In speech recognition, variations in the tempo of speech should, again, be treated as reparametrizations of the same underlying pattern, as should frequency variations.

In the Bayesian approach to the problem, one starts by assuming a space of probability distributions, $\{w\}$, within which the true distribution lies, and given a set of observations, f , one obtains a probabilistic description of the true distribution by using

$$\text{prob}(w|f) = \frac{\text{prob}(f|w) \cdot \text{prob}(w)}{\text{prob}(f)}; \quad (1)$$

in words, the probability of the distribution w being the true distribution, given f , is the probability of f given the distribution w , multiplied by the probability of w occurring in the space of all distributions, normalized by the probability of f occurring in *any* of the distributions in the set $\{w\}$. The *maximum likelihood* (ML) estimate of w is then that distribution w that maximizes $\text{prob}(f|w) \cdot \text{prob}(w)$. To go further in this direction, one needs to figure out what the *a priori* probability $\text{prob}(w)$ should be, and how to compute the ML estimate. Clearly, to avoid bias, one wants the space of distributions to be as general as possible, consistent with computability.

In one dimension, Bialek, Callan and Strong [2] have recently given an elegant formulation of this problem for continuous distributions in one dimension. They used Bayes' rule to write the probability of the probability distribution Q , given the data $\{x_i\}$, as

$$P[Q(x)|x_1, x_2, \dots, x_N] = \frac{P[x_1, x_2, \dots, x_N|Q(x)]P[Q(x)]}{P(x_1, x_2, \dots, x_N)} = \frac{Q(x_1)Q(x_2) \cdots Q(x_N)P[Q(x)]}{\int [dQ(x)]Q(x_1)Q(x_2) \cdots Q(x_N)P[Q(x)]}, \quad (2)$$

where the factors $Q(x_i)$ arise because each x_i is chosen independently from the distribution $Q(x)$, and $P[Q]$ encodes the *a priori* hypotheses about the form of Q . The optimal least-square estimate of Q , $Q_{\text{est}}(x, \{x_i\})$, is then

$$Q_{\text{est}}(x; \{x_i\}) = \frac{\langle Q(x)Q(x_1)Q(x_2) \cdots Q(x_N) \rangle^{(0)}}{\langle Q(x_1)Q(x_2) \cdots Q(x_N) \rangle^{(0)}}, \quad (3)$$

where $\langle \cdots \rangle^{(0)}$ denotes an expectation value with respect to the *a priori* distribution $P[Q(x)]$.

In this field-theoretic setting, Ref. [2] assumed that the prior distribution $P[Q]$ should penalize large gradients, so written in terms of an unconstrained variable $\phi \equiv \ln(\ell Q) \in (-\infty, +\infty)$, they assumed

$$P_\ell[\phi(x)] = \frac{1}{Z} \exp \left[-\frac{\ell}{2} \int dx (\partial_x \phi)^2 \right] \times \delta \left[1 - \frac{1}{\ell} \int dx e^{\phi(x)} \right], \quad (4)$$

with ℓ a *global* length parameter that they later fixed by means of other considerations. This form for the prior distribution is very simple, and quite minimal in terms of underlying assumptions, so is almost an ideal example of the Bayesian set-up.

There is, however, an important aspect in which this prior distribution does not measure up—a probability distribution is a density, and hence transforms in a specific manner under reparametrizations of the data, *e.g.* given a set of N observations x_i , we want to infer a distribution w_x that is related in a very simple way to the distribution w_{exp} inferred from $\exp(x_i)$:

$$w_x(z) = \exp(z) \cdot w_{\text{exp}}(\exp(z)). \quad (5)$$

This reparametrization covariance of the estimated distribution turns out to be a strong constraint on the distribution $\text{prob}(w)$, and on the actual computation of the estimated distribution. In recent work, I found a geometric modification of the approach of Ref. [2], which satisfies this desired reparametrization covariance [4]. It is the details of this solution, especially the corrections for finite sample size (which are obviously of paramount importance for practical applications), that are explained in the present paper.

This paper is organized as follows: Section 2 contains a short review of Ref. [4], to keep the presentation here relatively self-contained. In Section 3, the one-dimensional solution for finite sample sizes is given, along with a discussion of corrections arising from fluctuations about the saddlepoint distribution. In section 4, I present some conclusions, with comments on the solution of the inference problem in dimensions larger than 3.

II. REVIEW

Instead of the form (eq. 4) used in Ref. [2], I write $Q(x) \equiv \sqrt{h(x)} \exp(\phi(x))$, with $h(x)$ a metric in one dimension. Hence $\sqrt{h(x)}$ is a scalar density, and so is Q , which is crucial for a reparametrization covariant solution. An intuition for the rôle that h plays is that it is a local analogue of ℓ , determining binning intervals locally. Set

$$P[\phi, h] \equiv \frac{1}{Z} \exp \left[-\frac{\ell}{2} \int dx \sqrt{h(x)}^{-1} (\partial_x \phi)^2 \right] \delta \left[1 - \frac{1}{\ell} \int dx \sqrt{h(x)} e^{\phi(x)} \right], \quad (6)$$

The inverse of the metric is just $1/h(x)$ and the reparametrization invariant volume element is $\sqrt{h(x)} dx$, so $P[\phi, h]$ is reparametrization invariant. Now, we want to evaluate

$$\begin{aligned} & \langle Q(x_1) Q(x_2) \cdots Q(x_N) \rangle^{(0)} \\ &= \int D\phi \frac{Dh}{Dif f_+} P[\phi, h] \prod_{i=1}^N \sqrt{h} \exp[\phi(x_i)] \end{aligned} \quad (7)$$

$$= \frac{1}{Z} \int \frac{d\lambda}{2\pi} \int D\phi \frac{Dh}{Dif f_+} \exp[-S(\phi, h; \lambda)], \quad (8)$$

where

$$S(\phi, h; \lambda) = \frac{1}{2} \int dx \sqrt{h(x)}^{-1} (\partial_x \phi)^2 + i\lambda \int dx \sqrt{h(x)} e^{\phi(x)} - \sum_{i=1}^N [\phi(x_i) + \frac{1}{2} \ln h(x_i)] - i\lambda. \quad (9)$$

Notice that the integral over all metrics has been divided by the volume of the group of orientation preserving diffeomorphisms. In one dimension, this division eliminates all but one global degree of freedom from the metric. There is no operational way to distinguish between the factor $\sqrt{h(x_i)}$ and $\exp(\phi(x_i))$ —these must occur together in $Q(x_i)$. Taking the local symmetry into account, the number of local degrees of freedom is the same as in the approach of Bialek et al. [2].

The equations of motion that follow from varying S are

$$-\frac{1}{2}(\phi')^2 \frac{1}{h} + i\lambda \exp(\phi) - \sum_i \frac{1}{\sqrt{h(x_i)}} \delta(x - x_i) = 0, \quad (10)$$

$$-\left(\frac{1}{\sqrt{h}} \phi' \right)' + i\lambda \sqrt{h} \exp(\phi) - \sum_i \delta(x - x_i) = 0, \quad (11)$$

$$\int dx \sqrt{h} \exp(\phi) = 1, \quad (12)$$

where I have used primes to denote d/dx . $\delta(x)$ denotes the scalar density such that $\int dx \delta(x) = 1$. Introduce a variable

$$y(x) = \int^x \sqrt{h(s)} ds, \quad (13)$$

then it follows from eq. 10 and eq. 11 that

$$\frac{d^2\phi}{dy^2} = \frac{1}{2} \left(\frac{d\phi}{dy} \right)^2. \quad (14)$$

Note that the use of eq. 13 is limited by the non-singularity of \sqrt{h} .

It is now necessary to be careful about the limits of integration. This care in the boundary terms is unnecessary at $N = \infty$, in accord with the fact that any finite data set will not indicate the true limits of the probability distribution. Suppose that x ranges from x_- to x_+ , then integrating eq. 11 I find

$$N - i\lambda + \int_{x_-}^{x_+} dx \left(\frac{1}{\sqrt{h}} \phi' \right)' = 0, \quad (15)$$

so, using eq. 12, it follows that

$$\exp(-\phi) = (y_+ - y_-) \frac{(y - c)^2}{(y_- - c)(y_+ - c)} \quad (16)$$

where c is an arbitrary constant of integration. c is restricted only by $c > y_+ > y_-$, or $y_+ > y_- > c$, since $\exp(\phi)$ must be positive. This ‘cyclic’ constraint on y_+, y_-, c arises because the algebra of infinitesimal reparametrizations of the line has a subalgebra isomorphic to $sl(2, \mathbf{R})$.

In order to determine $y(x)$, which satisfies

$$\frac{(y_+ - c)(y - y_-)}{(y_+ - y_-)(y - c)} = \int_{x_-}^x \left[\frac{1}{N} \sum_i \delta(x' - x_i) \right] dx', \quad (17)$$

observe that the cross ratio on the left is projectively invariant, *i.e.*, invariant under transformations of the form

$$z \mapsto \frac{\alpha z + \beta}{\gamma z + \delta}, \quad (18)$$

with $\alpha, \beta, \gamma, \delta$ real, $\alpha\delta - \beta\gamma = 1$. This amounts to a three-parameter family of equivalent solutions $y(x)$ determined by the data. This projective invariance can be fixed by setting $c = \infty, y_+ = 1$ and $y_- = 0$, which implies $\phi = 0$.

Operationally, this means that the next measured data point will be observed with equal probability at any value of y in the interval $[0, 1]$. An important point to notice about this solution is that it is only valid at $N = \infty$. At finite N , the metric determined by this solution is not smooth, being a sum over δ functions, hence we must take more care in solving eq.’s 10,11,12 for finite N . Indeed, the action evaluated at this solution is infinite for finite N .

III. FINITE N CORRECTIONS

Obviously, for any practical application of this theory, one needs a systematic scheme for computing finite sample size corrections. To understand the theory at finite N , and to make contact with the higher-dimensional theory commented on in Section 4, we rewrite eq. 9 as

$$S = \frac{1}{2} \int dy (\partial_y \phi)^2 + i\lambda \left(\int dy e^\phi - 1 \right) + N \left[- \int dy \hat{P}_y(y) \phi(y) - \int dx \hat{P}_x \ln \sqrt{h(x)} \right], \quad (19)$$

where $y(x)$ is as in eq. 13, $\hat{P}_x(x) \equiv \frac{1}{N} \sum \delta(x - x_i)$, and $\hat{P}_y(y)y' \equiv \hat{P}_x(x)$. It is clear in this form that there are two independent functional variations of S , one with respect to $\phi(y)$, and the other with respect to $y(x)$:

$$-\partial_y^2 \phi + i\lambda \exp(\phi) - N \hat{P}_y(y) = 0, \quad (20)$$

$$\partial_x \left[\frac{\hat{P}_x(x)}{y'(x)} \right] = 0. \quad (21)$$

For N large, suppose $\hat{P}_x(x) = P(x) + \frac{1}{\sqrt{N}}\rho(x)$, with $P(x)$ a continuous density and $\rho(x) : \langle \rho(x)\rho(x') \rangle = P(x)\delta(x-x')$, following Ref. [2]. Now, a solution to eq.'s 20 and 21 at leading order in N is $\phi_0 = 0$, and $\sqrt{h}(x) \equiv \partial_x y_0(x) = P(x)$, as reviewed in Sect. 2. Indeed, it would seem from eq. 17 that this solution is exact, and no modification for finite N is necessary. This is *not* true, since for finite sample sizes eq. 17 defines a singular metric, and hence is inconsistent with our calculations which have assumed that the metric is non-singular implicitly.

We can, however, solve eq. 11 to get $\phi = \phi[\sqrt{h}, \hat{P}_x]$, and then use this solution in eq. 10 to determine h . We shall find that the leading correction at finite N is a correction to ϕ at $O(N^{-1/2})$, with no correction to $\sqrt{h} = P(x)$. This ordering of the computational steps parallels the calculation in [2], with the determination of h here the analogue of the determination of the length scale in [2]. What is different in the present approach, is that the formulæ derived here are reparametrization-covariant, *e.g.*, square roots of densities do not appear. To leading order, we find $i\lambda = N$, and

$$\exp(\phi_0) = \frac{P(x)}{\sqrt{h}}. \quad (22)$$

Notice that eq. 22 has the correct reparametrization covariance by construction, with a function equated to a ratio of densities.

At the next order, expanding eq. 11 about ϕ_0 , $i\lambda = N$, we find

$$\left[-\frac{1}{\sqrt{h}(x)} \partial_x \frac{1}{\sqrt{h}(x)} \partial_x + N e^{\phi_0} \right] \phi_1 \equiv (-\Delta_h + N e^{\phi_0}) \phi_1 = \left[\sqrt{N} \frac{\rho(x)}{\sqrt{h}(x)} + \Delta_h \phi_0 \right]. \quad (23)$$

Eq. 23 is easily solved to obtain

$$\phi_1(x) = \int dx' \sqrt{h}(x') \left[\sqrt{N} \frac{\rho(x')}{\sqrt{h}(x')} + \Delta_h \phi_0(x') \right] K(x, x'), \quad (24)$$

with the leading term in K given by

$$K(x, x') \approx \frac{1}{2\sqrt{N}} [e^{\phi_0}(x) e^{\phi_0}(x')]^{-1/4} \exp \left(-\sqrt{N} \int_{\min(x, x')}^{\max(x, x')} \sqrt{h} e^{\phi_0/2} dx'' \right). \quad (25)$$

Note the reparametrization invariance of eq. 24.

We can now insert eq.'s 22, 24 in eq. 9, and vary with respect to h to obtain $h = h(\hat{P}_x)$. This order of solving the saddlepoint equations is somewhat simpler than solving the equations obtained by independent variations—the results are, obviously, independent of this order of solution. We find

$$S_{red}[h] = -N \int dx \left[P + \frac{1}{\sqrt{N}} \rho \right] \ln P + \frac{1}{2} \int \left(NP \phi_1 [\ln \sqrt{h} - \ln P - \phi_1] - \sqrt{N} \rho [\ln \sqrt{h} + \phi_1 - \ln P] \right), \quad (26)$$

with $\phi_1 = \phi_1(h)$. Now, it is important to note that while S_{red} appears to include terms with powers of N that should be ignored at first glance, the variation of ϕ_1 with respect to h is $O(1)$ even though $\phi_1 = O(N^{-1/2})$. Thus, varying eq. 26 with respect to h , we find

$$\frac{1}{2} NP \frac{\delta \phi_1}{\delta \sqrt{h}} (\ln \sqrt{h} - \ln P) + O(\sqrt{N}) = 0$$

so to leading order in N , $\sqrt{h_0} = P$, $\phi_0 = 0$ so our old result is recovered. However, at the next order, we find *no* correction to h of $O(N^{-1/2})$. Thus the leading finite size correction is entirely encoded in eq. 24, with $\sqrt{h} = P$. At higher orders in $N^{-1/2}$, there are, of course, further corrections to ϕ , and to h .

Consider now the computation of the functional integral expanded about the saddlepoint solution found above. The division by the volume of the group of diffeomorphisms eliminates the integration over fluctuations in the metric, so we are left with computing the integral over the fluctuations in ϕ . This integration gives the operator determinant $\det^{-1/2}(-\Delta_h + N e^{\phi_1})$. This determinant can be computed in the large N limit by the standard van Vleck technique, explained in Ref. [5], for example. Thus

$$\det(-\Delta_{P^2} + N e^{\phi_1})^{-1/2} \propto \exp \left(-\frac{\sqrt{N}}{2} \int dx P(x) e^{\phi_1/2} \right), \quad (27)$$

which differs from the corresponding expression in Ref. [2]. Eq. 27 is reparametrization invariant. Putting everything together, eq. 8 is given by

$$\prod_{i=1}^N P(x_i) \exp \left(-\frac{\sqrt{N}}{2} \int dx P(x) e^{\phi_1/2} - \frac{1}{2} \int dx P(x)^{-1} (\partial_x \phi_1)^2 \right). \quad (28)$$

Eq. 28 can be used to evaluate the least-square estimate of the inferred distribution.

Finally, we compute the magnitude of the correction ϕ_1 :

$$\langle \phi_1(x)^2 \rangle_{\text{connected}} \sim \frac{1}{4\sqrt{N}}, \quad (29)$$

independent of x , implying *uniform* convergence to the solution reviewed in Sect. II. This uniform convergence is of great practical importance, since it implies that our estimated distribution will make sense even in regions where P is small. The uniformity can be directly traced to the reparametrization invariance of our computations, which enables the binning size to be adjusted locally enabling a more accurate determination of the distribution. As this uniformity is of great importance for practical applications, it is all the more gratifying that a re-analysis motivated by conceptual concerns regarding reparametrization covariance leads to a computational improvement.

IV. CONCLUSIONS

The statistical inference problem is of great interest in biophysics, for example, when x is a vector variable. Bialek, Callan and Strong [2] gave an account of a higher dimensional variant of their theory, but I want to add a few remarks here regarding a geometric version sketched in [4], taking into account reparametrization invariance. In Sect.'s II and III, we saw that a reparametrization-covariant Bayesian set-up could be obtained by introducing a metric, *and* a scalar field, dividing out by the group of diffeomorphisms. Since a metric in one dimension has only one global degree of freedom, this left only one local scalar degree of freedom, which is what we expect to need to fit the data, since the data determines a scalar density on the interval $[x_-, x_+]$.

In higher dimensions, the aim is still to determine a scalar density, but the number of local degrees of freedom in the metric in d dimensions is $d(d+1)/2$ and reparametrizations form a d parameter local symmetry group. It is clear then that a geometric formulation will require a reparametrization invariant constraint to reduce the number of degrees of freedom in the metric, to the required $d+1$ local degrees of freedom.

Reparametrization invariant constraints must be given by the vanishing of tensor quantities, so in this case we have to construct appropriate tensors from the metric. Investigating the constraints provided by the vanishing of various curvature tensors, since we expect only a scalar local degree of freedom to survive after imposing the constraint and dividing out by reparametrizations, we are naturally led to consider conformally Euclidean metrics, which can indeed be characterized by the vanishing of a tensor constructed from the curvature, the Weyl curvature, $W_{\alpha\beta\gamma}^\delta$. W vanishes identically in $d \leq 3$, so the geometric theory will only be valid for $d > 3$.

The proposed generalization of eq. 6 to higher dimensions is now given by

$$P[h_{\alpha\beta}] \propto \exp \left(- \int d^d x \sqrt{h} \mathcal{L} \right) \delta \left(1 - \int d^d x \sqrt{h(x)} \right) \prod_x \delta (W_{\alpha\beta\gamma}^\delta(x)) \quad (30)$$

where \mathcal{L} is a scalar constructed out of the metric $h_{\alpha\beta}$, constrained to be a metric of vanishing Weyl curvature, $W_{\alpha\beta\gamma}^\delta$, and $\sqrt{h} \equiv \det^{1/2}(h_{\alpha\beta})$ as usual.

A variety of issues, some technical, have to be addressed in order to define eq. 30 precisely: (i) The integration measure for integrating over metrics is notoriously nonlinear; (ii) Solving the constraint $W = 0$, and then extracting the volume of the group of diffeomorphisms will leave a Jacobian; (iii) What is an appropriate choice for \mathcal{L} , given the constraint $W = 0$?

Before tackling the technical issues, consider the logic of what we wish to do. Metrics with vanishing Weyl curvature are of the form

$$h_{\alpha\beta}(x) = \frac{\partial f^\delta}{\partial x^\alpha} \frac{\partial f^\epsilon}{\partial x^\beta} e^{2\phi(f(x))} \delta_{\delta\epsilon} \equiv J_\alpha^\delta J_\beta^\epsilon e^{2\phi(f(x))} \delta_{\delta\epsilon}. \quad (31)$$

For such metrics, $\sqrt{h} = \det(J) e^{d\phi}$. Now, following our steps in Sect. 2, we want to extremize (ignoring technicalities)

$$\hat{S} \equiv \int d^d x \det(J) e^{d\phi} \mathcal{L}[\phi(f(x))] + i\lambda \left(\int d^d x \det(J) e^{d\phi} - 1 \right) - N \int d^d x \hat{P}(x) [\ln \det J + \phi] , \quad (32)$$

where $N\hat{P}(x) \equiv \sum \delta(x - x_i)$. Now, $\hat{P}(x)$ is a density, so $d^d x \hat{P}(x) = d^d f \hat{P}(f)$, with $\hat{P}(f) \det(J) = \hat{P}(x)$. We can rewrite eq. 32 as

$$\hat{S} = \int d^d f e^{d\phi} \mathcal{L}[\phi(f)] + i\lambda \left(\int d^d f e^{d\phi} - 1 \right) - N \int d^d x \hat{P}(x) \ln \det J - N \int d^d f \hat{P}(f) \phi , \quad (33)$$

It follows that there are two independent functional variations, one with respect to $\phi(f)$, and the other with respect to $f(x)$, the latter giving

$$\frac{\partial}{\partial x^\alpha} \left[\hat{P}(x) \frac{\partial x^\alpha}{\partial f^\delta}(x) \right] = 0. \quad (34)$$

Eq. 34 is the analogue of eq. 21, and determines the coördinates f^α in terms of the coördinates x^α in which the data is presented, and is reparametrization-invariant because $P(x)$ is a density, not a function. Consider the special case $x = f$: eq. 34 is the requirement that the density P in the coördinates f is the *constant* density, exactly analogous to our result in one dimension, with ϕ a constant.

It is amusing to explicitly construct f^α coördinates which satisfy

$$\det \frac{\partial f^\alpha}{\partial x^\beta} = \hat{P}(x).$$

We construct these coördinates essentially iteratively as a one-dimensional distributions with parameters. If we could find f :

$$f^1(x^\alpha) = f^1(x^1), \quad f^2(x^\alpha) = f^2(x^1, x^2), \dots$$

then we would get

$$\det \frac{\partial f^\alpha}{\partial x^\beta} = \prod_\alpha \frac{\partial f^\alpha}{\partial x^\alpha},$$

so the problem is one of factoring \hat{P} appropriately. Define

$$\begin{aligned} P^1(x^1) &\equiv \int \hat{P}(x^1, x^2, x^3, \dots) \prod_{\alpha=2}^d dx^\alpha, \\ P^2(x^1, x^2) P^1(x^1) &\equiv \int \hat{P}(x^1, x^2, x^3, \dots) \prod_{\alpha=3}^d dx^\alpha, \\ &\vdots \\ P^d(x^\alpha) \prod_{\beta=1}^{d-1} P^\beta &\equiv \hat{P}(x^\alpha), \end{aligned} \quad (35)$$

then

$$\begin{aligned} f^1(x^1) &\equiv \int^{x^1} P^1(z) dz, \\ f^2(x^1, x^2) &\equiv \int^{x^2} P^2(x^1, z) dz, \\ &\vdots \end{aligned} \quad (36)$$

is an explicit set of coordinates in which the given density \hat{P} is constant. Notice that each f^α varies from 0 to 1, as befits integrals of one-dimensional normalized probability distributions, P^α .

As explained in Ref.'s [2,4], we are interested in extracting a continuum distribution from some set of binned data, as independent as possible of the binning. In renormalization group terminology, we are interested in the 'infrared' properties of the data, independent of the 'cutoff'. The terms most relevant in the infrared are precisely the terms with the fewest derivatives. Thus, since Einstein-Hilbert action is a total derivative for conformally Euclidean metrics, we consider $\mathcal{L} \equiv -\frac{1}{2}R\Delta^{-1}R$, where Δ is the Laplacian, and R is the Ricci scalar constructed from $h_{\alpha\beta}$. This Lagrange density is non-local for general metrics in dimensions larger than two, but is *local* for metrics with $W = 0$,

$$\int d^d x \sqrt{h} \mathcal{L} = \frac{(d-1)^2}{2} \int d^d f e^{(d-2)\phi(f)} \delta^{\alpha\delta} \frac{\partial \phi}{\partial f^\alpha} \frac{\partial \phi}{\partial f^\delta}. \quad (37)$$

The steps carried out in the previous section for determining the finite N corrections can be carried out in an exactly analogous fashion for this action for the leading correction, using the Laplacian in d dimensions in the f coordinates explicitly constructed above. At higher orders in N , the nonlinearity of \mathcal{L} leads to a distinctly different computation. I hope to address these issues, along with the technical issues mentioned above, elsewhere.

In conclusion, I have computed the finite sample size corrections to the geometric solution of the statistical inference problem I gave in earlier work. The uniform convergence of these corrections should be important for applications so it is of interest to test this theoretical construction in practice.

I am grateful to V. Balasubramanian and C. Callan for conversations. This work was supported in part by NSF grant PHY96-00258.

-
- [1] J. Earman, *Bayes or Bust?: A Critical Examination of Bayesian Confirmation Theory* (MIT Press, Cambridge, 1992); from a physics perspective, the inference problem has been discussed by V. Balasubramanian, "Statistical Inference, Occam's Razor and Statistical Mechanics on the Space of Probability Distributions", *Neural Computation* (in press), cond-mat/9601030; "Occam's Razor for Parametric Families and Prior on the Space of Probability Distributions", in *Maximum Entropy and Bayesian Methods*, K.M. Hanson and R.N. Silver (eds.), (Kluwer Academic, Amsterdam, 1996); "A Geometric Formulation of Occam's Razor for Inference of Parametric Distributions", adap-org/9601001; and references therein
 - [2] W. Bialek, C.G. Callan and S.P. Strong, *Phys. Rev. Lett.* **77** 4693(1996)
 - [3] C. Callan explained to me the biophysical applications in pattern recognition which motivated the investigation in Ref. [2]. See also D. Mumford, *Pattern Theory: A Unifying Perspective*, in *Proceedings, First European Congress of Mathematicians*, vol. 1, eds. A. Joseph, F. Mignot, B. Prum and R. Rentschler (Birkhäuser Verlag, Basel, 1992), and references therein.
 - [4] V. Periwal, *Phys. Rev. Lett.* **78** 4671(1997)
 - [5] S. Coleman, Ch. 7, *Aspects of symmetry* (Cambridge University Press, Cambridge, 1975)